

Literature Review on Cloud Computing Service Providers

Lee Jia Xian, Luqman Ariff Bin Noor Azhar, Adam Wafii Bin Azuar, Adriana Asyiqin Binti Md Adha, Kong Jia Rou

Universiti Teknologi Malaysia, Sultan Ibrahim Chancellery Building, Jalan Iman, 81310 Skudai, Johor, Malaysia

1. Introduction

Cloud computing is the availability of computer system resources on demand, especially data storage (cloud storage) and computing power, without the user's direct active management. The term is typically used to describe data centers accessible across the Internet to several users. Large clouds, prevalent today, also have functions spread from central servers over several locations

Cloud computing is the delivery of various services via the Internet. These resources include tools and applications like data storage, servers, databases, networking, and software. Cloud-based storage allows files to be saved to a remote database rather than storing them on a proprietary hard drive or a local storage unit. It has access to the data and the software programs to run it as long as an electronic device has access to the Internet.

2. Comparative Evaluation

In this comparative evaluation, we will base on service model, virtual machine (VM) instance type, storage, OS environment offered, security, performance and scalability, pricing model, auto-scaling and monitoring tools to compare the differences between Google Cloud Platform (GCP) and Microsoft Azure.

2.1. Service model (IaaS, PaaS, SaaS)

There are several types of service models offered by Google Cloud and Azure. Some of which are Infrastructure-as-a-service (IaaS), Platform-as-a-service (PaaS), and many more. For IaaS, the third-party provider hosts and maintains the core infrastructure such as hardware, software, servers and storage on behalf of the customers. PaaS on the other hand offers services such as a platform to the customers, allowing them to develop, manage and maintain the infrastructure. The difference between the two services are, the service provider will do most of the work for IaaS while the customers will do most of it for PaaS. Both Google Cloud and Azure offer IaaS and PaaS to their customers. For Google Cloud, it offers App Engine for their PaaS, which includes all the popular languages. It is also open and flexible for the customers as you can bring any library and framework to App Engine. For Azure, it offers App Service for their PaaS, which includes apps such as web apps, mobile apps, logic apps, and API apps.

2.2. Virtual Machine (VM) instance type offered

For Amazon EC2, there are various kinds of instance type offered. They dedicate certain resources of the computer, such as memory, CPU and instance storage. Amazon EC2 provides instance with a predictable and

constant amount of CPU capacity regardless the processors and memory. It also shares other resources of the host computer, such as the network and the disk subsystem, among instances. In the event that each instance on a host computer attempts to use these shared resources as possible, each gets an equivalent portion of that asset. In any case, when a resource is underused, an instance can consume a higher portion of that resource while it's accessible.

Meanwhile for Google Cloud Platform, they provide small boot persistent disk for each Compute Engine by default. They allow users to add additional storage options to their instances while applications are running on the instance. They also support declarative method for launching user application using containers. They offer logical packaging mechanisms for applications that can be absorbed from surrounding in which they can actually run.

Microsoft Azure has a variety of VMs with various pricing. A-series provide entry level workloads like development and test CPU performance and memory configuration. DC-series virtual machines protect the integrity of PC host data. This is in addition to the existing built-in encryption capabilities that protect data in Azure while it's at rest and in transit.

2.3.Storage

Azure storage services offer basic storage services such as Blob Storage for unstructured REST object data storage, Queue Storage for high-volume workloads, File Storage and Disk Storage. It also has a Data Lake Store, which is useful for big data applications. Azure storage services provide basic storage services like Blob Storage for unstructured REST object data storage, Queue Storage for high volume workloads, File Storage and Disk Storage. It features 3 SQL options: SQL Database, Database for MySQL and Database for PostgreSQL. There is also a Data Warehouse service, as well as Cosmos DB and Table Storage for NoSQL. Redis Cache is its memory service and the Server Stretch database is its hybrid storage service designed specifically for organizations that use Microsoft SQL Server within their own data centres. In contrast to AWS, Microsoft offers a real backup service, as well as site recovery and archive storage service.

Whereas within Google Storage, GCP has a growing menu of available storage services. Cloud Storage is its single object storage service, and it also has a persistent drive option. It offers a Transfer Appliance similar to AWS Snowball, as well as online transfer services. As far as databases are concerned, GCP has SQL Cloud SQL and a relational database called Cloud Spanner, designed to handle critical workloads. Two NoSQL options are also available: Cloud Bigtable and Cloud Datastore. It does not possess a backup and archive services.

2.4.OS environment offered

Both Azure and Google Cloud are compatible on all main Operating Systems such as Windows, Mac OS and Linux. The two services are offered to almost anyone regardless of their choice of Operating System which can be considered as a user-friendly service. Though for Mac and Linux users, they need to install certain software such as Microsoft Azure SDK Installer for Mac OS and the latest version of Node.js for Linux. Same goes for Google Cloud which requires an extra step on running it for Mac OS and Linux. All the instructions can be looked up online with all the steps that explains everything to the customers.

2.5.Security

We focus on the underlying technologies, controls, processes, and policies that combine to protect your cloud-based applications, data, and infrastructure when talking about cloud protection. Microsoft and Google are known for their strong dedication to delivering cloud protection at the highest standards. Compliance with your cloud platform is crucial as regulatory oversight of information from both governments and industry continues to

grow. Strict security procedures and processes are applied by both Google and Azure to ensure that they meet some of the toughest compliance criteria, including CSA STAR, GDPR, HIPPA, PCI-DSS, and a number of ISO specifications. To date, compliance with Azure is the highest of any cloud service, meeting 90+ levels of compliance spanning 50 global regions. Compliance with Google is also noteworthy, meeting 45 criteria for compliance. As for encryption of your files, regardless of whether it's in the cloud, is a vital necessity. Encoding the data means that it is almost difficult to decode without a decryption key should it be intercepted. Azure and Google Cloud support encryption within their cloud infrastructure, using 256-bit AES by default. They also give you the ability to monitor your own encryption keys and provide encryption at rest and in transit. Google refers to its service as the Cloud Key Management Service, while its Azure service is referred to as the Key Vault. Additional firewall-as-a-service products are provided by Azure, including its Azure Firewall, Azure Web Application Firewall, and the recently released cloud-native Azure Firewall Manager. Google has Cloud Identity and Access Management (IAM) and Microsoft has Azure Active Directory. Both providers offer an in-built IAM system. Microsoft has an edge of over 3,500 cybersecurity experts in terms of numbers, compared to Google's 550 cybersecurity experts. Through their respective Azure and Google Vulnerability Reward Programs, both providers also actively tap the expertise of the wider cybersecurity market.

2.6. Performance and scalability

In this first category of computing, we'll be concentrating on virtual machines (VMs). Comparison and contrast between Microsoft Azure and Google Cloud Platform installation. At a glance, each provider adopts a standardized approach to VMs, which form a fundamental part of any cloud environment, and will run almost every type of customer workload you can think of. Cloud giants come with different naming conventions for virtual machines. These are referred to as Azure virtual machines on Microsoft Azure and Compute Engine on Google Cloud Platform. The two providers also use different high-level terms and concepts.

For clarity of comparison, Google has mapped the differences for Azure and Compute Engine, as you can see in the table below:

Feature	Azure	Compute Engine
Virtual machines	Virtual machines	Virtual machine instances
Images	Image (both boot-disk-only and full machine)	Image (boot-disk-only)
Custom images	Generalized Azure VMs	Custom images
VM templates	Resource Manager templates	Instance templates

Automatic instance scaling	Azure Autoscale	Compute Engine autoscale
Supported VM import formats	VHD	RAW
Deployment locality	Regional (equivalent to Cloud Platform zones)	Zonal
Preemptible VMs	Yes	Yes
Incremental snapshots	Yes	Yes

2.7.Pricing model

One of the most challenging aspects of comparison when looking at cloud providers is pricing. There are so many variables, with each provider providing a distinctive pricing and billing approach. As for Microsoft Azure in order for an attempt to satisfy the specific needs of each buyer, the pricing choices are predominantly situational. Like AWS, Azure provides 750 hours of connecting Windows or Linux B1S to its main computing platform, Virtual Machines, every year. Also the prices are very flexible, so to find the best prices you might make a few surveys investing your money. When it comes down to pricing, the only aspect where Google aims to stand out from the other cloud services by making their pricing system a little less ambiguous and more customer friendly. They aim to beat most of cloud service providers' list prices and offer steep discounts and other rewards to win business. The free tier reward from Google includes one instance of F1-micro per month for up to one year. Google stands out for making it easier to work advanced tech into your budget in addition to their AI development. With this platform, you can get standard prices for basic services which, in their own way, are still revolutionary and unique. Finally, this is the platform for you if you are looking for an easy to use, budget-friendly service that shows promising growth potential.

2.8.Auto-scaling

Auto-scaling is a feature of cloud computing that enables cloud services, such as virtual machines (VM) and server capacities, to be automatically up or down, depending on the specific situation (eg. CPU utilization). Auto-scaling ensures that during demand spikes, new instances are increased smoothly and decreased during demand drops, allowing consistent performance for lower costs. Google Cloud Platform (GCP) and Microsoft Azure are the cloud providers that both offer Auto scaling features. The major differences of auto-scaling feature offered by GCP and Microsoft Azure are based on

There are Virtual Machine Scale Sets supported by Microsoft Azure. The user can build and manage a group of heterogeneous, load-balanced virtual machines with Azure Virtual Machine Scale Sets (VMs). Automatically increase or decrease the number of VMs in response to demand or based on schedule specified. Manage, configure, and update thousands of VMs centrally and provide the applications with greater availability and security. The scale sets provide key benefits, such as simple development and management of multiple VMs, high availability and application resilience, automatic scale of the application as resource demand shifts, and large-scale operation.

In Google Cloud Platform the auto-scaling works with managed instance groups (MIGs). Managed instance Groups (MIGs) provide auto-scaling capabilities that allow you to add or remove instances of a virtual machine (VM) from a MIG automatically based on load increases or decreases. Auto-scaling allows applications to gracefully manage traffic spikes and reduce costs where there is a lower demand for resources. When there is more load (scaling out, also referred to as scaling up), auto-scaling works by adding more VMs to your MIG and removing VMs when the need for VMs is reduced (scaling in or down).

2.9.Monitoring tools

As Microsoft Azure is a sophisticated cloud service, it can be difficult to understand and ensure the health and efficiency of the applications running on it. This necessitates the use of monitoring software from Azure. The Azure monitoring tool from Applications Managers helps to keep our services up and running by capturing, analyzing, and acting on telemetry from our cloud and on-premises environment. To maintain optimum efficiency, it provides deep insight into our applications and lets us recognize and address problems affecting them. So, users can use end-to-end monitoring tools with Azure monitoring tools to identify or diagnose problems in applications or infrastructure, manage Azure in virtual machine (VMs) or containers, detect bottlenecks, and collect data from a wide range of other tasks.

Since Google Cloud Platform is a public cloud infrastructure hosting a wide variety of resources, including software, storage, big data, machine learning (ML), creation of applications, and more. Google Cloud requires continuous monitoring to keep its overall health and performance intact. So, GCP provided three main types of monitoring tools. First, Google Cloud Storage Monitoring is used to proactively monitor critical areas. It helps in real-time tracking of storage buckets, quick troubleshooting of faults, intelligent warnings, and aggregated historical data analysis over a period. Second is the Google Cloud Filestore Monitoring. For applications that need a file system interface and a shared file system for data, Google Cloud Filestore is a managed file storage service. GCP Filestore monitoring helps to ensure the installation's health and availability, easily repair bugs, and deliver high performance. Lastly is the Google Kubernetes Engine also known as GKE is a cluster manager and orchestration system for running Docker containers in the cloud. Google Kubernetes Engine monitoring is the process of aggregating metrics and events from Kubernetes environment on GKE to understand the development activity of our application. It is important to monitor KPIs that define the performance of cloud.

3.Opinion

Each platform has their pros and cons no matter how close to perfect it is. For Amazon EC2 it is deemed as reliable as they offer high range of availability for each Amazon EC2 area. Replacements of instances can also be done easily. Their security services also prevent user from being exposed and remain private to the Internet. It is also flexible and cost saving. Their inexpensive pricing allow user to choose variety of plans and requirements. While they provide many benefits they also have limitation, such as not every region has Amazon EC2. There are limited number of instances per area. Limited information, they also charges user for immediate support. Google cloud has the lowest pricing among the three options and it is also available anywhere. Nowadays, almost everyone uses google and it is unexpectedly easy to use. Free users could use google cloud with limited storage but it is still quite impressive compared to the other options. The only downside of Google Cloud is they lack features compared to AWS and they provide small components. Moving on to Microsoft Azure, they have high availability and redundancy in data centers. Their main focus is on security, where they achieve multiple compliance certifications. The cons are user needs platform expertise to consult and teach them to ensure all moving parts work together efficiently.

In our humble opinion, we would rather choose Google Cloud as it is probably one of the money saving platform for students. We could purchase other options but considering the importance, Google Cloud is more suitable for students.

4.Conclusion

In terms of services, Azure provides a comprehensive package of services and storage features, but can have a steep learning curve, especially for users without experience in Microsoft technology. Google provides less functionality, but shines in storage pricing and ease of use. Google Cloud cannot yet compete with Azure's massive data center infrastructure, but reaches up for it with stronger support for container and Kubernetes use cases, and a smoother learning curve across all types of deployments. When it comes to pricing, Google Cloud rivals Microsoft Azure on pricing and offers more flexible pricing across virtually all cloud services. However, Azure offers a discount model that may be appealing to existing Microsoft clients.

References

1. Amazon.com. (2017). "Instance types - Amazon Elastic Compute Cloud." <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-types.html>
2. Ranger, Steve. "What Is Cloud Computing? Everything You Need to Know about the Cloud Explained." *ZDNet*, ZDNet, 13 Dec. 2018, <https://www.zdnet.com/article/what-is-cloud-computing-everything-you-need-to-know-about-the-cloud/>
3. Ranger, Steve. "What Is Cloud Computing? Everything You Need to Know about the Cloud Explained." *ZDNet*, ZDNet, 13 Dec. 2018, <https://www.e-zest.com/microsoft-azure-auto-scaling-implementation#:~:text=The%20auto%2Dscaling%20feature%20enables,by%20flaking%20unnecessary%20instances%20automatically>
4. Google Cloud. "Autoscaling Groups of Instances | Compute Engine Documentation." Google, Google, <https://cloud.google.com/compute/docs/autoscaler>
5. Bwren. "Azure Monitor Overview - Azure Monitor." *Azure Monitor | Microsoft Docs*, <https://docs.microsoft.com/en-us/azure/azure-monitor/overview>

